arXiv:cs/9905013v1 [cs.LG] 20 May 1999

Robust Combining of Disparate Classifiers through Order Statistics

Kagan Tumer
NASA Ames Research Center
MS 269-2, Moffett Field, CA, 94035-1000
kagan@ptolemy.arc.nasa.gov
Joydeep Ghosh
Department of Electrical and Computer Engineering,
University of Texas, Austin, TX 78712-1084
ghosh@pine.ece.utexas.edu

February 1, 2008

Abstract

Integrating the outputs of multiple classifiers via combiners or meta-learners has led to substantial improvements in several difficult pattern recognition problems. In the typical setting investigated till now, each classifier is trained on data taken or resampled from a common data set, or (almost) randomly selected subsets thereof, and thus experiences similar quality of training data. However, in certain situations where data is acquired and analyzed on-line at several geographically distributed locations, the quality of data may vary substantially, leading to large discrepancies in performance of individual classifiers. In this article we introduce and investigate a family of classifiers based on order statistics, for robust handling of such cases. Based on a mathematical modeling of how the decision boundaries are affected by order statistic combiners, we derive expressions for the reductions in error expected when such combiners are used. We show analytically that the selection of the median, the maximum and in general, the i^{th} order statistic improves classification performance. Furthermore, we introduce the trim and spread combiners, both based on linear combinations of the ordered classifier outputs, and show that they are quite beneficial in presence of outliers or uneven classifier performance. Experimental results on several public domain data sets corroborate these findings.

1 Introduction

Since different types of classifiers have different "inductive bias", one does not expect the generalization performance of two classifiers to be identical [22, 25] for difficult pattern recognition problems, even when they are both trained on the same data set. If only the "best" classifier is selected based on an *estimation* of the true generalization performance using a finite test set [60], valuable information contained in the results of the discarded classifiers may be lost. Such potential loss of information can be avoided if the outputs of all available classifiers are used in the final classification decision. This

concept has received a great deal of attention recently, and many methods for combining classifier outputs have been proposed [23, 27, 29, 42, 53]. Furthermore, diversity among classifiers has been actively promoted, by strategies such as bagging [8], arcing [9, 19, 20], boosting [18, 17, 45, 51, 52], and correlation control [2, 59], as a prelude to combining.

Approaches to pooling classifiers can be separated into two main categories: (i) simple combiners, e.g., voting [4, 12], Bayesian based weighted product rule [31], or averaging [41, 58], and, (ii) meta-learners, such as arbitration [11] or stacking [7, 61]. The simple combining methods are best suited for problems where the individual classifiers perform the same task, and have comparable success. However, such combiners are more susceptible to outliers and to unevenly performing classifiers. In the second category, either sets of combining rules, or full fledged classifiers acting on the outputs of the individual classifiers, are constructed [1, 30, 61]. This type of combining is more general, but is vulnerable to all the problems associated with the added learning (e.g., overparameterizing, lengthy training time).

An implicit assumption in most combining schemes is that each classifier sees the same training data or resampled versions of the same data. If the individual classifiers are then appropriately chosen and trained properly, their performances will be (relatively) comparable in any region of the problem space. So gains from combining are derived from the diversity [32, 40] among classifiers rather that by compensating for weak members of the pool. However, in real life, there are situations where individual classifiers may not have access to the same data. Such conditions arise in certain data mining, sensor fusion and electrical logging (oil services) problems where there are large variabilities in the data which is acquired locally and needs to be processed in (near) real time at geographically separated places [13]. These conditions create a pool of classifiers that may have significant variations in their overall performance. Moreover, they may lead to conditions where individual classifiers have similar *average* performance, but substantially different performance over different parts of the input space.

In such cases, combining is still desirable, but neither simple combiners nor meta-learners are particularly well-suited for the type of problems that arise. For example, the simplicity of averaging the classifier outputs is appealing, but the prospect of one poor classifier corrupting the combiner makes this a risky choice. Weighted averaging of classifier outputs appears to provide some flexibility [28, 37]. Unfortunately, the weights are still assigned on a per classifier basis rather than a per sample or per class basis. If a classifier is accurate only in certain areas of the input space, this scheme fails to take advantage of the variable accuracy of the classifier in question. Using a meta learner that provides different weights for different patterns can potentially solve this problem, but at a considerable cost. In particular, the off-line training of a meta-learner using substantial amount of data outputted by geographically distributed classifiers, may not be feasible. In addition to providing robustness, the order statistic combiners presented in this work also aim at bridging the gap between simplicity and generality by allowing the flexible selection of classifiers without the associated cost of training meta-classifiers.

Section 2 summarizes the relationship between classifier errors and decision boundaries and provides the necessary background for mathematically analyzing order statistic combiners [58]. Section 3 introduces simple order statistic combiners. Based on these concepts, in Section 4 we propose two powerful combiners, *trim* and *spread*, and derive the amount of error reduction associated with each. In Section 5 we present the performance of order statistic combiners on Proben1/UCI benchmarks [43]. Section 6

discusses the implications of using linear combinations of order statistics as a strategy for pooling the outputs of individual classifiers.

2 Error Characterization in a Single Classifier

In this section we summarize the approach and results of [58][1], that quantify the effect of inaccuracies in estimating *a posterior* class probabilities on the classification error for a *single* classifier. This background is needed to characterize and understand the impact of order statistics combiners, as described in Sections 3 and 4.

It is well known that, given *one-of-L* desired outputs and sufficient training samples reflecting the class priors, the outputs of certain classifiers trained to minimize a mean square or cross-entropy error criteria, approximate the *a posteriori* probability densities of the corresponding classes [47, 49]. Based on this result, one can model the ith output of the mth such classifier as:

$$f_i^m(x) = p_i(x) + \epsilon_i^m(x), \tag{1}$$

where $p_i(x)$ is the true posterior for ith class on input x, and $\epsilon_i^m(x)$ is the error of the mth classifier in estimating that posterior.

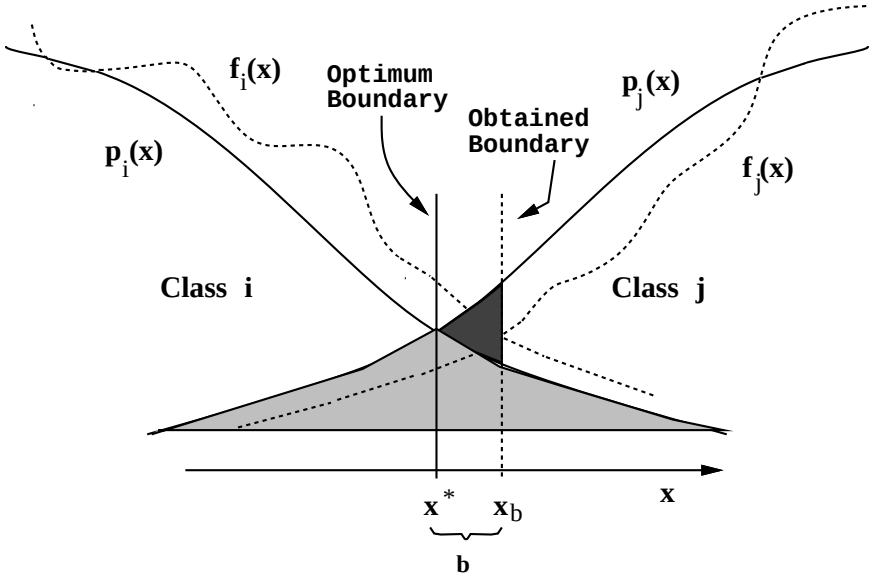

Figure 1: Error regions associated with approximating the *a posteriori* probabilities [58].

Now, let us decompose the error into two parts: $\epsilon_i^m(x) = \beta_i^m + \eta_i^m(x)$. The first component does not vary with the input, and provides an offset, or systematic error for each class. The second component gives the variability from that systematic error, for each x in each class, and has zero mean and variance $\sigma^2_{\eta_i^m(x)}$. These two components of the error are similar to the bias and variance decomposition for a quadratic loss function given in [22], although they are at the individual input level. We will therefore refer to classifiers as "biased" and "unbiased" implying $\beta_k^m \neq 0$ for some k, m, and

[1]This and other related papers can be downloaded from URL *http://www.lans.ece.utexas.edu.*

$\beta_k^m = 0$, $\forall k, m$, respectively. Let b^m denote the offset between the ideal class boundary, x^* (based on $p_i(x) = p_j(x)$) and the realized boundary, x_b^m (based on $f_i^m(x) = f_j^m(x)$), as shown in Figure 1 [58]. This boundary offset ($b^m = x_b^m - x^*$) has mean and variance given respectively by:

$$\beta^m = \frac{\beta_i^m - \beta_j^m}{s}, \tag{2}$$

and

$$\sigma_{b^m}^2 = \frac{\sigma_{\eta_i^m(x)}^2 + \sigma_{\eta_j^m(x)}^2}{s^2}, \tag{3}$$

where $s = p_j'(x^*) - p_i'(x^*)$ as introduced in [58].

Let us further denote the probability density function of this boundary offset by $f_b(x)$. The expected model error associated with the selection of a particular classifier m, can then be expressed as:

$$E_{model}^m = \int_{-\infty}^{\infty} A(b) f_b(b) db, \tag{4}$$

where $A(b) = \int_{x^*}^{x^*+b} (p_j(x) - p_i(x))\, dx$ is the error due to the selection of a particular decision boundary. In general, it is not possible to obtain the density function for the boundary offset without making assumptions on the distributions of the errors. However, a first order approximation, derived in [58], leads to:

$$E_{model}^m = \int_{-\infty}^{\infty} \frac{1}{2} b^2 s f_b(b) db. \tag{5}$$

Let us define the first and second moments of the boundary offset as follows:

$$\mathcal{M}_1 = \int_{-\infty}^{\infty} x f_b(x) dx \quad and \quad \mathcal{M}_2 = \int_{-\infty}^{\infty} x^2 f_b(x) dx.$$

If the individual classifiers are unbiased, the offset b^m of a single classifier has $\mathcal{M}_1 = 0$ and $\mathcal{M}_2 = \sigma_{b^m}^2$, leading to:

$$E_{model}^m = \frac{s\, \mathcal{M}_2}{2} = \frac{s\sigma_{b^m}^2}{2}. \tag{6}$$

Now, if the classifiers are biased, the variance of b is left unchanged (given by Equation 3), but the mean becomes $\beta = \frac{\beta_i - \beta_j}{s}$. In other words, we have $\mathcal{M}_1 = \beta^m$ and $\sigma_{b^m}^2 = \mathcal{M}_2 - {\mathcal{M}_1}^2$, leading to the following model error:

$$E_{model}^m(\beta) = \frac{s\mathcal{M}_2}{2} = \frac{s}{2} (\sigma_{b^m}^2 + (\beta^m)^2). \tag{7}$$

To emphasize the distinction between biased and unbiased classifiers, the model error will be given as a function of β for biased classifiers. A more detailed derivation of class boundaries and error regions is presented in [58]. For analyzing the error regions after combining and comparing them to the single classifier case, one needs to determine how the first and second moments of the boundary distributions are affected by combining. The following sections focus on obtaining those values for various combiners.

3 Combining Multiple Classifiers through Order Statistics

3.1 Basic Concepts

In this section, we briefly discuss some basic concepts and properties of order statistics. Let X be a random variable with probability density function $f_X(\cdot)$, and cumulative distribution function $F_X(\cdot)$. Let $(X_1, X_2, \cdots, X_N)$ be a random sample drawn from this distribution. Now, let us arrange them in non-decreasing order, providing:

$$X_{1:N} \leq X_{2:N} \leq \cdots \leq X_{N:N}.$$

The ith order statistic denoted by $X_{i:N}$, is the ith value in this progression. The cumulative distribution function for the smallest and largest order statistic can be obtained by noting that:

$$F_{X_{N:N}}(x) = P(X_{N:N} \leq x) = \Pi_{i=1}^{N} P(X_{i:N} \leq x) = [F_X(x)]^N$$

and:

$$\begin{aligned} F_{X_{1:N}}(x) &= P(X_{1:N} \leq x) = 1 - P(X_{1:N} \geq x) = 1 - \Pi_{i=1}^{N} P(X_{i:N} \geq x) \\ &= 1 - (1 - \Pi_{i=1}^{N} P(X_{i:N} \leq x) = 1 - [1 - F_X(x)]^N \end{aligned}$$

The corresponding probability density functions can be obtained from these equations. In general, for the ith order statistic, the cumulative distribution function gives the probability that exactly i of the chosen X's are less than or equal to x. The probability density function of $X_{i:N}$ is then given by [14]:

$$f_{X_{i:N}}(x) = \frac{N!}{(i-1)!\,(N-i)!} \left[F_X(x)\right]^{i-1} \left[1 - F_X(x)\right]^{N-i} f_X(x)\ . \tag{8}$$

This general form however, cannot always be computed in closed form. Therefore, obtaining the expected value of a function of x using Equation 8 is not always possible. However, the first two moments of the density function are widely available for a variety of distributions [3]. These moments can be used to compute the expected values of certain specific functions, e.g., polynomials of order less than two.

3.2 Combining Unbiased Classifiers through Order Statistics

Now, let us turn our attention to order statistics (OS) combiners. For a given input x, let the network outputs of each of the N classifiers for each class i be ordered in the following manner:

$$f_i^{1:N}(x) \leq f_i^{2:N}(x) \leq \cdots \leq f_i^{N:N}(x).$$

Then one constructs the kth order statistic combiner, by selecting the kth ranked output for each class ($f_i^{k:N}(x)$), as representing its posterior [57].

In particular, *max*, *med* and *min* combiners are defined as follows:

$$f_i^{max}(x) = f_i^{N:N}(x), \tag{9}$$

$$f_i^{med}(x) = \begin{cases} \frac{f_i^{\frac{N}{2}:N}(x) + f_i^{\frac{N}{2}+1:N}(x)}{2} & \text{if } N \text{ is even} \\ f_i^{\frac{N+1}{2}:N}(x) & \text{if } N \text{ is odd,} \end{cases} \tag{10}$$

$$f_i^{min}(x) = f_i^{1:N}(x). \tag{11}$$

These three combiners are relevant because they represent important qualitative interpretations of the output space. Selecting the maximum combiner is equivalent to selecting the class with the highest posterior. Indeed, since the network outputs approximate the class *a posteriori* distributions, selecting the maximum reduces to selecting the classifier that is the most "certain" of its decision. The drawback of this method however is that it can be compromised by a single classifier that repeatedly provides high values. The selection of the minimum combiner follows a similar logic, but focuses on classes that are unlikely to be correct, rather than on the correct class. Thus, this combiner eliminates less likely classes by basing the decision on the lowest value for a given class. This combiner suffers from the same ills as the *max* combiner. However, it is less dependent on a single error, since it performs a min-max operation, rather than a max-max[2]. The median classifier on the other hand considers the most "typical" representation of each class. For highly noisy data, this combiner is more desirable than either the *min* or *max* combiners since the decision is not compromised as much by a single large error.

The analysis that follows does not depend on the particular order statistic chosen. Therefore, we will denote all OS combiners by $f_k^{os}(x)$ and derive the model error, E_{model}^{os}. The network output provided by $f_k^{os}(x)$ is given by:

$$f_k^{os}(x) = p_k(x) + \epsilon_k^{os}(x)\,, \tag{12}$$

Let us first investigate the zero-bias case ($\beta_k = 0\,, \forall k$), where we get $\epsilon_k^{os}(x) = \eta_k^{os}(x)$. Proceeding as in Section 2, the boundary b^{os} is shown to be:

$$b^{os} \;=\; \frac{\eta_i^{os}(x_b) - \eta_j^{os}(x_b)}{s}. \tag{13}$$

For *i.i.d.* η_k's, the first two moments will be identical for each class. Moreover, taking the order statistic will shift the mean of both η_i^{os} and η_j^{os} by the same amount, leaving the mean of the difference unaffected. Therefore, b^{os} will have zero mean, and variance:

$$\sigma_{b^{os}}^2 \;=\; \frac{2\,\sigma_{\eta_k^{os}}^2}{s^2} \;=\; \frac{2\,\alpha\sigma_{\eta_k^m}^2}{s^2} \;=\; \alpha\sigma_{b^m}^2, \tag{14}$$

where α is a reduction factor that depends on the order statistic and on the distribution of b. For most distributions, α can be found in tabulated form [3]. For example, Table 1 provides α values for all order statistic combiners, up to 10 classifiers, for a Gaussian distribution [3, 50]. (Because this distribution is symmetric, the α values of l and k where $l + k = N + 1$ are identical, and listed in parenthesis).

Returning to the error calculation, we have: $\mathcal{M}_1^{os} = 0$, and $\mathcal{M}_2^{os} = \sigma_{b^{os}}^2$, providing:

$$E_{model}^{os} = \frac{s\mathcal{M}_2^{os}}{2} = \frac{s\sigma_{b^{os}}^2}{2} = \frac{s\alpha\sigma_{b^m}^2}{2} = \alpha\; E_{model}^{m}. \tag{15}$$

Equation 15 shows that the reduction in the error due to using the OS combiner instead of the mth classifier is directly related to the reduction in the variance of the boundary offset b. Since the means and variances of order statistics for a variety of distributions are widely available in tabular form, the reductions can be readily quantified.

[2]Recall that the pattern is ultimately assigned to the class with the highest combined output.

Table 1: Reduction factors α for the Gaussian Distribution, based on [50].

N	k	α	N	k	α	N	k	α
1	1	1.00	6	2 (5)	.280		1 (9)	.357
2	1 (2)	.682		3 (4)	.246		2 (8)	.226
3	1 (3)	.560		1 (7)	.392	9	3 (7)	.186
	2	.449	7	2 (6)	.257		4 (6)	.171
4	1 (4)	.492		3 (5)	.220		5	.166
	2 (3)	.360		4	.210		1 (10)	.344
	1 (5)	.448		1 (8)	.373		2 (9)	.215
5	2 (4)	.312	8	2 (7)	.239	10	3 (8)	.175
	3	.287		3 (6)	.201		4 (7)	.158
6	1 (6)	.416		4 (5)	.187		5 (6)	.151

3.3 Combining Biased Classifiers through Order Statistics

In this section, we analyze the error regions for biased classifiers. Let us return our attention to b^{os}. First, note that the error terms can no longer be studied separately, since in general $(a+b)^{os} \neq a^{os} + b^{os}$. We will therefore need to specify the mean and variance of the result of each operation[3]. Equation 13 becomes:

$$b^{os} = \frac{(\beta_i + \eta_i(x_b))^{os} - (\beta_j + \eta_j(x_b))^{os}}{s}. \tag{16}$$

Let $\bar{\beta_k} = \frac{1}{N} \sum_{m=1}^{N} \beta_k^m$ be the mean of classifier biases. Since η_k^m's have zero-mean, $\beta_k + \eta_k(x_b)$ has first moment $\bar{\beta_k}$ and variance $\sigma^2_{\eta_k^m} + \sigma^2_{\beta_k^m}$, with $\sigma^2_{\beta_k^m} = E[(\beta_k^m)^2] - \bar{\beta_k}^2$, where $[\cdot]$ denotes the expected value operator.

Taking a specific order statistic of this expression will modify both moments. The first moment is given by $\bar{\beta_k} + \mu^{os}$, where μ^{os} is a shift which depends on the order statistic chosen, but not on the class. Then, the first moment of b^{os} is given by:

$$\frac{(\bar{\beta_i} + \mu^{os}) - (\bar{\beta_j} + \mu^{os})}{s} = \frac{\bar{\beta_i} - \bar{\beta_j}}{s} = \bar{\beta}. \tag{17}$$

Note that the bias term represents an "average bias" since the contributions due to the order statistic are removed. Therefore, reductions in bias cannot be obtained from a table similar to Table 1.

Now, let us turn our attention to the variance. Since $\beta_k^m + \eta_k^m(x_b)$ has variance $\sigma^2_{\eta_k^m} + \sigma^2_{\beta_k^m}$, it follows that $(\beta_k + \eta_k(x_b))^{os}$ has variance $\sigma^2_{\eta_k^{os}} = \alpha(\sigma^2_{\eta_k^m} + \sigma^2_{\beta_k^m})$, where α is the factor discussed in Section 3.2. Therefore, the variance of b^{os} is given by:

$$\begin{aligned} \sigma^2_{b^{os}} &= \frac{\sigma^2_{\eta_i^{os}} + \sigma^2_{\eta_j^{os}}}{s^2} = \frac{2\,\alpha\sigma^2_{\eta_i^m}}{s^2} + \frac{\alpha(\sigma^2_{\beta_i^m} + \sigma^2_{\beta_j^m})}{s^2} \\ &= \alpha(\sigma^2_{b^m} + \sigma^2_{\beta^m}), \end{aligned} \tag{18}$$

where $\sigma^2_{\beta^m} = \frac{\sigma^2_{\beta_i^m} + \sigma^2_{\beta_j^m}}{s^2}$ is the variance introduced by the systematic errors of different classifiers.

[3]Since the exact distribution parameters of b^{os} are not known, we use the sample mean and the sample variance.

We have now obtained the first and second moments of b^{os}, and can compute the model error. Namely, we have $\mathcal{M}_1^{os} = \bar{\beta}$ and $\sigma^2_{b^{os}} = \mathcal{M}_2^{os} - (\mathcal{M}_1^{os})^2$, leading to:

$$E^{os}_{model}(\beta) = \frac{s}{2}\mathcal{M}_2^{os} = \frac{s}{2}(\sigma^2_{b^{os}} + \bar{\beta}^2) \quad (19)$$

$$= \frac{s}{2}(\alpha(\sigma^2_{b^m} + \sigma^2_{\beta^m}) + \bar{\beta}^2). \quad (20)$$

The reduction in the error is more difficult to assess in this case. By writing the error as:

$$E^{os}_{model}(\beta) = \alpha\frac{s}{2}(\sigma_b^2 + (\beta^m)^2) + \frac{s}{2}(\alpha\sigma^2_\beta + \bar{\beta}^2 - \alpha(\beta^m)^2),$$

we get:

$$E^{os}_{model}(\beta) = \alpha\, E^m_{model}(\beta) + \frac{s}{2}(\alpha\sigma^2_\beta + \bar{\beta}^2 - \alpha(\beta^m)^2). \quad (21)$$

Analyzing the error reduction in the general case requires knowledge about the bias introduced by each classifier. Unlike regression problems where the bias and variance contributions to the error are additive and well-understood, in classification problems their interaction is more complex [21]. Indeed it has been observed that ensemble methods do more than simply reduce the variance [52].

Based on these observations and Equation 21, let us analyze extreme cases. For example, if each classifier has the same bias, σ^2_β is reduced to zero and $\bar{\beta} = \beta^m$. In this case the error reduction can be expressed as:

$$E^{os}_{model}(\beta) = \frac{s}{2}(\alpha\sigma_b^2 + (\beta^m)^2 = \alpha E^m_{model}(\beta) + \frac{s(1-\alpha)}{2}(\beta^m)^2,$$

where α balances the two contributions to the error. A small value for α will reduce the first component of the error (mainly variance), while leaving the second term untouched. The net effect will be very similar to results obtained for regression problems. In this case, it is important to reduce classifier bias before combining (e.g., by using an overparametrized model).

If on the other hand, the biases produce a zero mean variable, we obtain $\bar{\beta} = 0$. In this case, the model error becomes:

$$E^{os}_{model}(\beta) = \alpha\, E^m_{model}(\beta) + \frac{s\,\alpha}{2}(\sigma^2_{\beta^m} - (\beta^m)^2)$$

and the error reduction will be significant if the second term is small or negative. In fact, if the variation among the biases is small relative to their magnitude, the error will be reduced more than in the unbiased cases. If however, the variation is large compared to the magnitude, the error reduction will be minimal. Furthermore, if α is large and the biases are small and highly varied, it is possible for this combiner to do worse than the individual classifiers, which is a danger not present for regression problems. This observation very closely parallels results reported in [21].

4 Linear Combining of Ordered Classifier Outputs

In the previous section, we derived error reductions when the class posteriors are directly estimated through the ordered classifier outputs. Since simple averaging has also been shown to provide benefits, in this section, we investigate the combinations of averaging and order statistics for pooling classifier outputs.

4.1 Spread Combiner

The first linear combination of ordered classifier outputs we study focuses on extrema. As discussed in Section 3.2, the maximum and minimum of a set of classifier outputs carry specific meanings. Indeed, the maximum can be viewed as the class for which there is the most evidence. Similarly, the minimum deletes classes with little evidence. In order to avoid a single classifier from having too large of an impact on the eventual output, these two values can be averaged to yield the *spread* combiner. This combiner strikes a balance between the positive and negative evidence, leading to a more robust combiner than either of them.

4.1.1 Spread Combiner for Unbiased Classifiers:

For a classifier without bias, the spread combiner is formally defined as:

$$f_i^{spr}(x) = \frac{1}{2}\left(f_i^{1:N}(x) \;+\; f_i^{N:N}(x)\right) \;=\; p(c_i|x) \;+\; \eta_i^{spr}(x)\,, \tag{22}$$

where:

$$\eta_i^{spr}(x) = \frac{1}{2}\;\left(\eta_i^{1:N}(x) \;+\; \eta_i^{N:N}(x)\right)\;.$$

The variance of $\eta_i^{spr}(x)$ is given by:

$$\sigma^2_{\eta_i^{spr}} \;\;=\;\; \frac{1}{4}\sigma^2_{\eta_i^{1:N}(x)} \;+\; \frac{1}{4}\,\sigma^2_{\eta_i^{N:N}(x)} \;+\; \frac{1}{2}cov(\eta_i^{1:N}(x), \eta_i^{N:N}(x)). \tag{23}$$

where $cov(\cdot,\cdot)$ represents the covariance between two variables (even when the η_i's are independent, ordering introduces correlations). Note that because of the ordering, the variances in the first two terms of Equation 23 can be expressed in terms of the individual classifier variances. Furthermore, the covariance between two order statistics can also be determined in tabulated form for given distributions. Table 2 provides these values for a Gaussian distribution based on [50]. This expression can be further simplified for symmetric distributions where $\sigma^2_{\eta^{1:N}} = \sigma^2_{\eta^{N:N}}$ (e.g., Gaussian noise model) and leads to:

$$\sigma^2_{\eta_i^{spr}} \;\;=\;\; \frac{1}{2}\left(\alpha_{1:N} + B_{1,N:N}\right)\sigma^2_{\eta_i(x)}, \tag{24}$$

where $\alpha_{m:N}$ is the variance of the mth ordered sample and $B_{m,l:N}$ is the covariance between the mth and lth ordered samples, given that the initial samples had unit variance [50]. Because this is a symmetric distribution, the β values are also symmetric (e.g., $\beta_{1,2:5} = \beta_{4,5:5}$).

Then, using Equation 3, the variance of the boundary offset b^{spr} can be calculated:

$$\begin{aligned}\sigma^2_{b^{spr}} \;\;&=\;\; \frac{\sigma^2_{{\eta_i}^{spr}} \;+\; \sigma^2_{{\eta_j}^{spr}}}{s^2} \\ &=\;\; \frac{1}{2}\left(\alpha_{1:N} + B_{1,N:N}\right)\sigma_b^2.\end{aligned} \tag{25}$$

Finally, through Equation 6, we can obtain the reduction in the model error due to the spread combiner:

$$\frac{E^{spr}_{model}}{E_{model}} \;=\; \frac{\alpha_{1:N} + B_{1,N:N}}{2}\;. \tag{26}$$

Table 2: Some Reduction Factors B for the Gaussian Distribution, based on [50].

N	k,l	B	N	k,l	B	N	k,l	B	N	k,l	B
2	1,2	.318		2,3	.189		1,4	.095		1,6	.059
3	1,2	.276	6	2,4	.140		1,5	.075		1,7	.049
	1,3	.165		2,5	.106		1,6	.060		1,8	.040
	1,2	.246		3,4	.183		1,7	.048		1,9	.031
4	1,3	.158		1,2	.196		1,8	.037		2,3	.154
	1,4	.105		1,3	.132		2,3	.163		2,4	.117
	2,3	.236		1,4	.099	8	2,4	.123		2,5	.093
	1,2	.224		1,5	.077		2,5	.098		2,6	.077
	1,3	.148		1,6	.060		2,6	.079	9	2,7	.063
5	1,4	.106	7	1,7	.045		2,7	.063		2,8	.052
	1,5	.074		2,3	.175		3,4	.152		3,4	.142
	2,3	.208		2,4	.131		3,5	.121		3,5	.114
	2,4	.150		2,5	.102		3,6	.098		3,6	.093
	1,2	.209		2,6	.080		4,5	.149		3,7	.077
	1,3	.139		3,4	.166		1,2	.178		4,5	.137
6	1,4	.102		3,5	.130		1,3	.121		4,6	.113
	1,5	.077		1,2	.186	9	1,4	.091			
	1,6	.056	8	1,3	.126		1,5	.073			

Based on Equation 26 and Tables 1 and 2, Table 3 displays the error reductions provided by the spread combiner for a Gaussian noise model (for comparison purposes, the error reduction for the *min* and *max* combiners is also provided. Note that for the Gaussian distribution, the error reduction of *min* is equal to that of *max*.).

Table 3: Error Reduction Factors for the Spread, *min* and *max* Combiners with Gaussian Noise Model.

N	*spread*	*min* or *max*
2	.500	.682
3	.362	.560
4	.299	.492
5	.261	.448
6	.236	.416
7	.219	.392
8	.205	.373
9	.194	.357
10	.186	.344

4.1.2 Spread Combiner for Biased Classifiers:

Now, if the classifier biases are non-zero, the spread combiner's output is given by:

$$f_i^{spr}(x) = \frac{1}{2}\left(f_i^{1:N}(x) \; + \; f_i^{N:N}(x)\right) \; = \; p(c_i|x) \; + \; (\eta_i(x) + \beta_i)^{spr} \; . \tag{27}$$

In that case, the boundary offset is given by:

$$b^{spr} = \frac{(\beta_i + \eta_i(x_b))^{spr} - (\beta_j + \eta_j(x_b))^{spr}}{s}, \tag{28}$$

which after expanding each term and regrouping can be expressed as:

$$\begin{aligned} b^{spr} &= \frac{(\beta_i + \eta_i(x_b))^{1:N} - (\beta_j + \eta_j(x_b))^{1:N}}{2s} \\ &+ \frac{(\beta_i + \eta_i(x_b))^{N:N} - (\beta_j + \eta_j(x_b))^{N:N}}{2s}. \end{aligned} \tag{29}$$

The first moment of b^{spr} can be obtained by analyzing each term of Equation 29. In fact, the offset introduced by the first and nth order statistic for classes i and j will cancel each other out, leaving only the average bias between the min and max components of the error (as in Equation 17), given by $\beta^{spr} = \frac{\beta_i^{1:N} - \beta_j^{1:N} + \beta_i^{N:N} - \beta_j^{N:N}}{s}$.

The variance of b^{spr} needs to be derived from Equation 29. Proceeding as in Equation 18, the variance of the spread combiner can be expressed as:

$$\sigma^2_{b^{spr}} = (\frac{1}{4}\alpha_{1:N} + \frac{1}{4}\alpha_{N:N} + \frac{1}{2}B_{1,N:N})(\sigma^2_{b^m} + \sigma^2_{\beta^m}). \tag{30}$$

For a symmetric distribution (where $\alpha_{1:N} = \alpha_{N:N}$), we obtain the following error:

$$\begin{aligned} E^{spr}_{model}(\beta) &= \frac{s}{2}\mathcal{M}_2 = \frac{s}{2}(\sigma^2_{b^{spr}} + {\mathcal{M}_1}^2) \\ &= \frac{s}{2}\left(\frac{1}{2}\alpha_{1:N} + \frac{1}{2}B_{1,N:N})(\sigma^2_{b^m} + \sigma^2_{\beta^m}) + (\beta^{spr})^2\right) \\ &= \frac{1}{2}(\alpha_{1:N} + B_{1,N:N})E_{model}(\beta) + \\ &\quad \frac{s}{4}(\alpha_{1:N} + B_{1,N:N})(\sigma^2_{\beta^m} - (\beta^m)^2) + \frac{s}{2}(\beta^{spr})^2, \end{aligned} \tag{31}$$

which is very similar to Equation 21, where the value of α for a single order statistic is now replaced by $\frac{\alpha_{1:N} + B_{1,N:N}}{2}$, since the mean of the first and nth order statistic is used in the posterior estimate.

4.2 Trimmed Means

Instead of actively using the extreme values as was the case with the spread combiner, one can base the posterior estimate around the median values. However, instead of selecting one classifier output as was done for f^{med}, one can use multiple classifiers whose outputs are "typical." In this scheme, only a certain fraction of all available classifiers are used *for a given* pattern. The main advantage of this method over weighted averaging is that the set of classifiers which contribute to the combiner vary from pattern to pattern. Furthermore, they do not need to be determined externally, but are a function of the current pattern and the classifier responses to that pattern.

4.2.1 Trimmed Mean Combiner for Unbiased Classifiers:

Let us formally define the trimmed mean combiner ($\beta_k = 0, \forall k$) as follows:

$$f_i^{trim}(x) = \frac{1}{N_2 - N_1 + 1} \sum_{m=N_1}^{N_2} f_i^{m:N}(x) = p(c_i|x) + \eta_i^{trim}(x), \tag{32}$$

where:

$$\eta_i^{trim}(x) = \frac{1}{N_2 - N_1 + 1} \sum_{m=N_1}^{N_2} \eta_i^m(x) \; .$$

The variance of $\eta_i^{trim}(x)$ is given by:

$$\begin{aligned} \sigma^2_{\eta_i^{trim}} &= \frac{1}{(N_2 - N_1 + 1)^2} \sum_{l=N_1}^{N_2} \sum_{m=N_1}^{N_2} cov(\eta_i^{m:N}(x), \eta_i^{l:N}(x)) \\ &= \frac{1}{(N_2 - N_1 + 1)^2} \left(\sum_{m=N_1}^{N_2} \sigma^2_{\eta_i^{m:N}(x)} + \sum_{m=N_1}^{N_2} \sum_{l>m}^{N_2} 2 \; cov(\eta_i^{m:N}(x), \eta_i^{l:N}(x)) \right) \; . \end{aligned} \quad (33)$$

Again, using the factors in Tables 1 and 2, Equation 33 can be further simplified. Note that because the Gaussian distribution is symmetric, the covariance between the kth and lth ordered samples is the same as that between the $N+1-k$th and $N+1-l$th ordered samples. Therefore, Equation 33 leads to:

$$\begin{aligned} \sigma^2_{\eta_i^{trim}} &= \frac{1}{(N_2 - N_1 + 1)^2} \sum_{m=N_1}^{N_2} \alpha_{m:N} \;\; \sigma^2_{\eta_i(x)} \\ &+ \frac{2}{(N_2 - N_1 + 1)^2} \sum_{m=N_1}^{N_2} \sum_{l>m} B_{m,l:N} \;\; \sigma^2_{\eta_i(x)} \; , \end{aligned} \quad (34)$$

where $\alpha_{m:N}$ is the variance of the mth ordered sample and $B_{m,l:N}$ is the covariance between the mth and lth ordered samples, given that the initial samples had unit variance [50]. Using the theory highlighted in Section 2, and Equation 34, we obtain the following model error reduction:

$$\frac{E^{trim}_{model}}{E_{model}} = \frac{1}{(N_2 - N_1 + 1)^2} \left(\sum_{m=N_1}^{N_2} \alpha_{m:N} + 2 \sum_{m=N_1}^{N_2} \sum_{l>m} B_{m,l:N} \right) \; . \quad (35)$$

Based on Equation 35 and Tables 1 and 2, we have generated a sample *trim* combiner reduction table. Because there are many possibilities for N_1 and N_2, a table that exhaustively provides all reduction values is not practical. In this sample table we have selected $N_1 = 2$ and $N_2 = N - 1$, that is, averaging after the lowest and highest values have been removed. For comparison purposes the reduction factors of the averaging combiner for N and $N-2$ classifiers are also provided (for i.i.d. classifiers the reduction factors are 1/N as derived in [58]; similar results were obtained for regression problems [42]). As these numbers demonstrate, although $N-2$ classifiers are used in the trim combiner, *selectively* weeding out undesirable classifiers provides reduction factors significantly better than simply averaging $N-2$ arbitrary classifiers. The *trim* combiner provides reduction factors comparable the the N classifier *ave* combiner without being susceptible to corruption by one particularly faulty classifier.

4.2.2 Trimmed mean Combiner for Biased Classifiers:

Now, if the classifier biases are non-zero, the trimmed mean combiner's output is given by:

$$f_i^{trim}(x) = \frac{1}{N_2 - N_1 + 1} \sum_{m=N_1}^{N_2} f_i^{m:N}(x) = p(c_i|x) + (\eta_i(x) + \beta_i)^{trim} \; . \quad (36)$$

Table 4: Error Reduction Factors for Trim and two corresponding *ave* Combiners with Gaussian Noise Model.

N	*ave* (for N)	*trim* (for $N_1 = 2$; $N_2 = N - 1$)	*ave* (for $N - 2$)
3	.333	.449	1.00
4	.250	.298	.500
5	.200	.227	.333
6	.167	.184	.250
7	.143	.155	.200
8	.125	.134	.167
9	.111	.113	.143

In that case the boundary offset is given by:

$$b^{trim} = \frac{(\beta_i + \eta_i(x_b))^{trim} - (\beta_j + \eta_j(x_b))^{trim}}{s}. \tag{37}$$

The first moment of b^{trim} can be obtained from a manner similar to that of the spread combiner. Indeed, each mean offset introduced by a specific order statistic for class i will be offset by the one introduced for class j. Only the trimmed mean of the biases will remain, giving the first moment of b^{trim}:

$$\beta^{trim} = \frac{1}{N_2 - N_1 + 1} \sum_{m=N_1}^{N_2} \frac{\beta_i^{m:N} - \beta_j^{m:N}}{s}. \tag{38}$$

In deriving the variance of b^{trim}, we follow the same steps as in Sections 3.3 and 4.1.1. The resulting boundary variance is similar to Equation 18, but the since the reduction is due to the linear combination of multiple ordered outputs, α is replaced by $\mathcal{A}$, where:

$$\mathcal{A} = \frac{1}{(N_2 - N_1 + 1)^2} \left(\sum_{m=N_1}^{N_2} \alpha_{m:N} \; + \; 2 \sum_{m=N_1}^{N_2} \sum_{l>m} B_{m,l:N} \right). \tag{39}$$

The model error reduction in this case is given by:

$$\begin{aligned} E_{model}^{trim}(\beta) &= \frac{s}{2}\mathcal{M}_2 \;=\; \frac{s}{2}(\sigma^2_{b^{trim}} \; + \; {\mathcal{M}_1}^2) \\ &= \frac{s}{2}\left(\mathcal{A} \;\; (\sigma^2_{b^m} + \sigma^2_{\beta^m}) + (\beta^{spr})^2\right) \\ &= \mathcal{A} \;\; E_{model}(\beta) \;\; + \;\; \frac{s}{2}(\mathcal{A} \;\; (\sigma^2_{\beta^m} - (\beta^m)^2) + (\beta^{spr})^2) \; . \end{aligned} \tag{40}$$

Once again we need to look at the interaction between the two parts of the error reduction. The first term provides the error reduction compared to the model error of an individual classifier. The smaller $\mathcal{A}$ is, the more error reduction there will be. In the second term, on the other hand, a small value for $\mathcal{A}$ is only useful if the variability in the individual biases is higher than the biases themselves ($\sigma^2_{\beta^m} > (\beta^m)^2$).

5 Experimental Results

The order statistics-based combining methods proposed in this article are tailored for situations where:

1. individual classifier performance is uneven and class dependent;
2. it is not possible (insufficient data, high amount of noise) to fine tune the individual classifiers without using computationally expensive methods.

Such situations occur, for example, in electrical logging while drilling for oil, where data from certain well sites almost completely misses out on portions of the problem space, and in imaging from airborne platforms where the classifiers receive inputs from different satellites and/or different types of sensors (e.g., thermal, optical, SAR). While we have seen such data from Schlumberger, Austin, and NASA, Houston, unfortunately the data sets are not standard or public domain. So, in this article we restrict ourselves to public domain datasets and simulate such variability by using "early stopping" i.e., prematurely terminating the training of the individual classifiers[4]. Thus combining results are first reported for the case where only half the classifiers are finely tuned. This procedure produces an artificially created *quality* variation in the pool of classifiers.

For the experiments reported below, we used a multi-layer perceptron (MLP) with a single hidden layer, whose weights were randomly initialized for each run. All classification results reported in this article are *test set error rates averaged over 20 runs, along with the* 95% *confidence intervals*. Several types of simple combiners such as averaging, weighted averaging, voting, median, products, weighted products (Bayesian), using Dempster-Schafer theory of evidence, and entropy-based averaging, have been proposed in the literature. However, on a wide variety of data sets, it has been observed that simple averaging usually provides results comparable to any of these techniques (and, surprisingly, often better than most of them) [26, 59]. For this reason, in this study, we use the average combiner as a *representative* of simple combiners, for comparison purposes.

The first two data sets (Tables 5 and 7) are based on underwater sonar signals. From the original sonar signals of four different underwater objects (porpoise sound, cracking ice and two different whale sounds), two feature sets are extracted [24]:

- WOC: a 25-dimensional feature set, consisting of Gabor wavelet coefficients, temporal descriptors and spectral measurements; and,
- RDO: a 24-dimensional feature set, consisting of reflection coefficients based on both short and long time windows, and temporal descriptors.

For both feature sets, an MLP with 50 hidden units was used. These data sets are available at URL *http://www.lans.ece.utexas.edu*. Further details about this 4-class problem can be found in [24, 59].

The next six data sets (Tables 6 and 8) were selected from the Proben1/UCI benchmarks [43]. The Proben1 benchmarks are particular training, validation and test splits of the UCI data sets which are available from URL http://www.ics.uci.edu/~mlearn/MLRepository.html. The results presented in this article are based on the first training, validation and test partition discussed in [43], where half the data is used for training, and a quarter each for validation and testing purposes. Briefly these data sets, and the corresponding single layer feed-forward neural network architectures are[5]:

[4]In all the experiments reported here, "high variability" among classifiers refers to classifiers being trained *exactly half as long* as the "fine tuned" classifiers.

[5]After deciding on a single hidden layered architecture, the number of hidden units was determined experimentally.

- Cancer: a 9-dimensional, 2-class data set based on breast cancer data [34], with 699 patterns; an MLP with 10 hidden units;
- Card: a 51-dimensional, 2-class data set based on credit approval decision [44], with 690 patterns; an MLP with 20 hidden units;
- Diabetes: an 8-dimensional data set with two classes based on personal data from 768 Pima Indians obtained from the National institute of Diabetes and Digestive and Kidney Diseases [54]; an MLP with 10 hidden units;
- Gene: a 120-dimensional data set with two classes, based on the detection of splice junctions in DNA sequences [39], with 3175 patterns; an MLP with 20 hidden units;
- Glass: a 9-dimensional, 6-class data set based on the chemical analysis of glass splinters, with 214 patterns; an MLP with 15 hidden units; and,
- Soybean: an 82-dimensional, 19-class problem [38] with 683 patterns; an MLP with 40 hidden units.

Table 5: Combining Results in the Presence of High Variability in Individual Classifier Performance for the Sonar Data (% misclassified ± 95% confidence interval).

Data	N	Ave	Max	Min	Spread	Trim (N_1-N_2)
RDO	4	$11.57 \pm .22$	$11.94 \pm .25$	$11.52 \pm .40$	$11.04 \pm .19$	$11.34 \pm .28$ (3-4)
13.32 ± 1.66	8	$11.64 \pm .18$	$11.47 \pm .22$	$11.29 \pm .27$	$11.51 \pm .18$	$12.30 \pm .17$ (4-5)
WOC	4	$8.80 \pm .18$	$7.84 \pm .20$	$9.31 \pm .24$	$8.54 \pm .12$	$8.43 \pm .26$ (3-4)
12.07 ± 2.23	8	$8.82 \pm .17$	$7.68 \pm .23$	$8.91 \pm .13$	$8.24 \pm .22$	$7.81 \pm .16$ (7-8)

Table 6: Combining Results in the Presence of High Variability in Individual Classifier Performance for the Proben1/UCI Benchmarks (% misclassified ± 95% confidence interval).

Data	N	Ave	Max	Min	Spread	Trim (N_1-N_2)
Cancer	4	$1.38 \pm .13$	$1.38 \pm .13$	$1.38 \pm .13$	$1.38 \pm .13$	$1.32 \pm .13$ (2-3)
$1.49 \pm .39$	8	$1.32 \pm .12$	$1.44 \pm .14$	$1.44 \pm .14$	$1.44 \pm .14$	$1.32 \pm .12$ (2-6)
Card	4	$13.60 \pm .22$	$13.37 \pm .22$	$13.49 \pm .21$	$13.37 \pm .22$	$13.60 \pm .15$ (3-4)
$14.33 \pm .36$	8	$13.66 \pm .19$	$13.08 \pm .14$	$13.02 \pm .14$	$12.97 \pm .12$	$13.20 \pm .18$ (7-8)
Diabetes	4	$25.26 \pm .37$	$25.00 \pm .46$	$25.00 \pm .42$	$25.00 \pm .42$	$25.26 \pm .37$ (3-4)
26.09 ± 1.27	8	$24.84 \pm .36$	$25.05 \pm .33$	$25.05 \pm .33$	$25.05 \pm .33$	$24.84 \pm .30$ (6-8)
Gene	4	$12.90 \pm .23$	$12.90 \pm .26$	$12.94 \pm .25$	$12.66 \pm .21$	$12.67 \pm .22$ (3-4)
$15.01 \pm .78$	8	$12.89 \pm .22$	$12.76 \pm .24$	$12.41 \pm .10$	$12.43 \pm .22$	$12.56 \pm .20$ (7-8)
Glass	4	$33.77 \pm .27$	$40.19 \pm .72$	$33.21 \pm .44$	$33.21 \pm .44$	$33.77 \pm .27$ (2-3)
$42.78 \pm .75$	8	$33.96 \pm .06$	$39.43 \pm .27$	$33.77 \pm .27$	$33.40 \pm .41$	$33.77 \pm .27$ (1-6)
Soybean	4	$7.76 \pm .11$	$7.94 \pm .14$	$12.88 \pm .39$	$7.71 \pm .15$	$7.82 \pm .18$ (3-4)
10.71 ± 1.69	8	$7.65 \pm .00$	$7.82 \pm .13$	$13.41 \pm .53$	$7.71 \pm .15$	$7.65 \pm .00$ (4-8)

Tables 5 and 6 present the combining results for the Proben1 benchmarks and the underwater acoustic data sets respectively, when the individual classifier performance

was highly variable. The misclassification percentage for individual classifiers are reported in the first column. For the trimmed mean combiner, we also provide N_1 and N_2, the upper and lower cutting points in the ordered average used in Equation 32, obtained through the validation set.

On the Sonar data, the results indicate that when the individual classifier performance is highly variable, order statistics-based combiners (particularly the *spread* combiner) provide better classification results than simple combiners. This performance improvement is obtained without sacrificing the simplicity of the combiner. On the UCI/Proben1 benchmarks, the order statistics based combiners provide better classification performance on three of the six sets studied (no statistically significant differences were detected among the various combiners in the remaining data sets). One important thing to note, however, is that in all eight data sets studied, the order statistics based combiners performed *at least as well* as the simple combiner, implying that no risk is taken by using this method.

A close inspection of these results reveals that using either the *max* or *min* combiner can provide better classification rates than *ave*, but it is difficult to determine which of the two will be more successful given a data set. A validation set may be used to select one over the other, but in that case, potentially precious training data is used solely for determining which combiner to use. The use of the *spread* combiner removes this dilemma by consistently providing results that are comparable to, or better than, the best of the *max-min* duo. It is important to note that the *min* combiner performs poorly on the Soybean data. Because this data set has 19 outputs, the posterior estimates of unlikely classes become extremely small and highly inaccurate. Basing decisions on such spurious values compromises the combiner's performance. Notice, however, that the *spread* combiner is not adversely affected by this phenomenon.

Table 7: Combining Results with Fine-Tuned Classifiers for the Sonar Data (% misclassified $\pm$ 95% confidence interval).

Data	N	Ave	Max	Min	Spread	Trim (N_1-N_2)
RDO	4	$9.26 \pm .32$	$9.67 \pm .20$	$9.45 \pm .19$	$9.33 \pm .20$	$9.28 \pm .28$ (2-3)
$9.95 \pm .36$	8	$8.94 \pm .06$	$9.62 \pm .16$	$9.36 \pm .15$	$9.48 \pm .18$	$8.92 \pm .10$ (1-6)
WOC	4	$7.05 \pm .12$	$7.31 \pm .15$	$7.44 \pm .17$	$7.31 \pm .16$	$7.05 \pm .16$ (2-3)
$7.47 \pm .21$	8	$7.17 \pm .08$	$7.19 \pm .12$	$7.41 \pm .16$	$7.22 \pm .07$	$7.07 \pm .10$ (2-6)

When there is ample data, and all the classifiers are finely tuned (i.e., a validation set is used to determine the stopping time that yields the best generalization performance), simple combiners are expected to be adequate. However, it is not always possible to determine whether all conditions that lead to such an ideal situation are satisfied. Therefore, it is important to know whether the trimmed mean and spread combiners presented in this article perform worse than simple combiners under such conditions. To that end, we have combined finely tuned feed forward neural networks using the methods proposed in this article and compared the results with the traditional averaging method. In this new set of experiments, all the conditions favor the averaging combiner (i.e., all possible difficulties for the average combiner have been removed). The results displayed in Tables 7 and 8 indicate that, even under such circumstances, both the *spread* and *trim* combiners provide results that are comparable to those obtained by the

ave combiner. Furthermore, even under such conditions, the order statistics combiners provide statistically significant improvements on two data sets.

Table 8: Combining Results with Fine-Tuned Classifiers for the Proben1/UCI Benchmarks (% misclassified $\pm$ 95% confidence interval).

Data	N	Ave	Max	Min	Spread	Trim (N_1-N_2)
Cancer	4	$0.69 \pm .11$	$0.69 \pm .11$	$0.69 \pm .11$	$0.69 \pm .11$	$0.69 \pm .11$ (2-3)
$.69 \pm .11$	8	$0.69 \pm .11$	$0.57 \pm .01$	$0.57 \pm .01$	$0.57 \pm .01$	$0.57 \pm .11$ (7-8)
Card	4	$13.14 \pm .23$	$12.91 \pm .11$	$13.02 \pm .23$	$12.91 \pm .11$	$13.14 \pm .23$ (2-3)
$13.87 \pm .36$	8	$13.14 \pm .23$	$12.79 \pm .01$	$12.79 \pm .01$	$12.79 \pm .01$	$12.80 \pm .01$ (7-8)
Diabetes	4	$23.33 \pm .29$	$23.23 \pm .30$	$23.33 \pm .24$	$23.23 \pm .30$	$23.33 \pm .29$ (3-4)
$23.52 \pm .35$	8	$22.92 \pm .23$	$23.23 \pm .34$	$23.12 \pm .34$	$23.23 \pm .34$	$22.92 \pm .23$ (4-8)
Gene	4	$12.41 \pm .21$	$12.46 \pm .24$	$12.51 \pm .18$	$12.41 \pm .17$	$12.41 \pm .12$ (3-4)
$13.49 \pm .21$	8	$12.26 \pm .14$	$12.46 \pm .18$	$12.16 \pm .08$	$12.11 \pm .19$	$12.16 \pm .09$ (1-6)
Glass	4	$32.08 \pm .01$	$32.45 \pm .36$	$32.08 \pm .01$	$32.08 \pm .01$	$32.08 \pm .01$ (3-6)
$32.26 \pm .27$	8	$32.08 \pm .01$	$32.08 \pm .01$	$32.08 \pm .01$	$32.08 \pm .01$	$32.08 \pm .01$ (3-6)
Soybean	4	$7.06 \pm .00$	$7.18 \pm .11$	$8.12 \pm .77$	$7.06 \pm .00$	$7.06 \pm .00$ (3-6)
$7.36 \pm .43$	8	$7.06 \pm .00$	$7.18 \pm .05$	$9.06 \pm .82$	$7.06 \pm .00$	$7.06 \pm .00$ (3-6)

6 Conclusion

In this article we present and analyze combiners based on order statistics. These combiners blend the simplicity of averaging with the generality of meta-learners. They are particularly effective if there are significant variations among component classifiers in at least some parts of the joint input-output space. Variations can arise when the individual training sets cannot be considered as *random* samples from a common universal data set. Examples of such cases include real-time data acquisition and classification from geographically distributed sources or data mining problems with large databases, where random subsampling is computationally expensive and practical methods lead to non-random subsamples [6]. Furthermore, The robustness of order statistics combiners is also helpful when certain individual classifiers experience catastrophic failures (e.g., due to faulty sensors).

The analytical framework provided in this paper quantifies the reductions in error achieved when an order statistics based ensemble is used. It also shows that the two methods for linear combination of order statistics introduced in this paper provide more reliable estimates of the true posteriors than any of the individual order statistic combiners.

The experimental results of Section 5 indicate that when there is high variability among the classifiers, the order statistics-based combiners significantly outperform simple combiners, whereas in the absence of such variability these combiners perform no worse. Thus the family of order statistic combiners is able to extract an appropriate amount of information from the individual classifier outputs without requiring tuning additional parameters as in meta-learners, and without being substantially affected by outliers.

A future endeavor, which will be helpful for this work as well as for the study of classification based on very large datasets in general, is to obtain a suite of public domain datasets which are intrinsically partitioned into segments with varying quality. Though such situations sometimes occur in practice (for example in oil logging data [10] and mortgage scoring [36]; both data sets proprietary), they are not represented in the standard, venerable databases such as UCI, ELENA and Statlog typically used by the academic community. Perhaps the recent CRoss-Industry Standard Process for Data Mining (CRISP-DM) initiative will provide a satisfactory solution to this problem in the near future.

Acknowledgements: This research was supported in part by AFOSR contract F49620-93-1-0307, ARO contracts DAAH04-94-G0417 and DAAH04-95-10494, and NSF grant ECS 9307632.

References

[1] K. Al-Ghoneim and B. V. K. Vijaya Kumar. Learning ranks with neural networks (Invited paper). In *Applications and Science of Artificial Neural Networks, Proceedings of the SPIE*, volume 2492, pages 446–464, April 1995.

[2] K. M. Ali and M. J. Pazzani. On the link between error correlation and error reduction in decision tree ensembles. Technical Report 95-38, Department of Information and Computer Science, University of California, Irvine, 1995.

[3] B.C. Arnold, N. Balakrishnan, and H.N. Nagaraja. *A First Course in Order Statistics*. Wiley, New York, 1992.

[4] R. Battiti and A. M. Colla. Democracy in neural nets: Voting schemes for classification. *Neural Networks*, 7(4):691–709, 1994.

[5] W. G. Baxt. Improving the accuracy of an artificial neural network using multiple differently trained networks. *Neural Computation*, 4:772–780, 1992.

[6] P.S. Bradley and U. M. Fayyad. Refining initial points for K--means clustering. In *Proceedings of the International Conference on Machine Learning (ICML-98)*, pages 91–99, July 1998.

[7] L. Breiman. Stacked regression. Technical Report 367, Department of Statistics, University of California, Berkeley, 1993.

[8] L. Breiman. Bagging predictors. Technical Report 421, Department of Statistics, University of California, Berkeley, 1994.

[9] L. Breiman. Bias, variance and arcing classifiers. Technical Report 460, Department of Statistics, University of California, Berkeley, 1996.

[10] S. Chakravarthy. Private communication. Western Atlas, Houston, 1997.

[11] P. Chan and S. Stolfo. On the accuracy of meta-learning for scalable data mining. *Journal of Intelligent Integration of Information*, 1997. (to appear).

[12] P.K. Chan and S.J. Stolfo. A comparative evaluation of voting and meta-learning on partitioned data. In *Proceedings of the Twelfth International Machine Learning Conference*, pages 90–98, Tahoe City, CA, 1995. Morgan Kaufmann.

[13] B. Dasarathy. *Decision Fusion*. IEEE CS Press, Los Alamitos, CA, 1994.

[14] H. A. David. *Order Statistics*. Wiley, New York, 1970.

[15] T. G. Dietterich and E. B. Kong. Error-correcting output coding corrects bias and variance. In *Proceedings of the Twelfth International Conference on Machine Learning*, pages 313–321, Tahoe City, CA, 1995. Morgan Kaufmann.

[16] T. G. Dietterich and E. B. Kong. Machine learning bias, statistical bias, and statistical variance of decision tree algorithms. Technical report, Department of Computer Science, Oregon State University, 1995.

[17] H. Drucker and C. Cortes. Boosting decision trees. In D. S. Touretzky, M. C. Mozer, and M. E. Hasselmo, editors, *Advances in Neural Information Processing Systems-8*, pages 479–485. M.I.T. Press, 1996.

[18] H. Drucker, C. Cortes, L. D. Jackel, Y. LeCun, and V. Vapnik. Boosting and other ensemble methods. *Neural Computation*, 6(6):1289–1301, 1994.

[19] Y. Freund and R. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. In *Proceedings of the Second European Conference on Computational Learning Theory*, pages 23–37. Springer Verlag, March 1995.

[20] Y. Freund and R. Schapire. Experiments with a new boosting algorithm. In *Proceedings of the Thirteenth International Conference on Machine Learning*, pages 148–156. Morgan Kaufmann, 1996.

[21] J. H. Friedman. On bias, variance, 0/1-loss, and the curse-of-dimensionality. 1997. (from: http://www-stat.stanford.edu∼jhf/#reports).

[22] S. Geman, E. Bienenstock, and R. Doursat. Neural networks and the bias/variance dilemma. *Neural Computation*, 4(1):1–58, 1992.

[23] J. Ghosh, S. Beck, and C.C. Chu. Evidence combination techniques for robust classification of short-duration oceanic signals. In *SPIE Conf. on Adaptive and Learning Systems, SPIE Proc. Vol. 1706*, pages 266–276, Orlando, Fl., April 1992.

[24] J. Ghosh, L. Deuser, and S. Beck. A neural network based hybrid system for detection, characterization and classification of short-duration oceanic signals. *IEEE Journal of Ocean Engineering*, 17(4):351–363, October 1992.

[25] J. Ghosh and K. Tumer. Structural adaptation and generalization in supervised feedforward networks. *Journal of Artificial Neural Networks*, 1(4):431–458, 1994.

[26] J. Ghosh, K. Tumer, S. Beck, and L. Deuser. Integration of local and global neural classifiers for passive sonar signals. In *Proceedings of the International Simulation Technology Conferenc*, pages 539–545, Houston, TX, 1992.

[27] L. K. Hansen and P. Salamon. Neural network ensembles. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(10):993–1000, 1990.

[28] S. Hashem and B. Schmeiser. Approximating a function and its derivatives using MSE-optimal linear combinations of trained feedforward neural networks. In *Proceedings of the Joint Conference on Neural Networks*, volume 87, pages I:617–620, New Jersey, 1993.

[29] T. K. Ho, J. J. Hull, and S. N. Srihari. Decision combination in multiple classifier systems. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(1):66–76, 1994.

[30] Robert Jacobs. Method for combining experts' probability assessments. *Neural Computation*, 7(5):867–888, 1995.

[31] J. Kittler, M. Hatef, R.P.W. Duin, and J. Matas. On combining classifiers. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(3):226–239, 1998.

[32] A. Krogh and J. Vedelsby. Neural network ensembles, cross validation and active learning. In G. Tesauro, D. S. Touretzky, and T. K. Leen, editors, *Advances in Neural Information Processing Systems-7*, pages 231–238. M.I.T. Press, 1995.

[33] M. LeBlanc and R. Tibshirani. Combining estimates in regression and classification. Technical report, University of Toronto, 1993.

[34] O. L. Mangasarian, R. Setiono, and W. H. Wolberg. Pattern recognition via linear programming: Theory and application to medical diagnosis. In Thomas F. Coleman and Yuying Li, editors, *Large-Scale Numerical Optimization*, pages 22–30. SIAM Publications, 1990.

[35] D. Margineantu and T. G. Dietterich. Pruning adaptive boosting. In *Proceedings of the Fourteenth International Conference on Machine Learning*, pages 211–218. Morgan Kaufmann, 1997.

[36] C. J. Merz. Private communication. HNC Software, San Diego, 1998.

[37] C.J. Merz and M.J. Pazzani. Combining neural network regression estimates with regularized linear weights. In M. C. Mozer, M. I. Jordan, and T. Petsche, editors, *Advances in Neural Information Processing Systems-9*, pages 564–570. M.I.T. Press, 1997.

[38] R.S. Michalski and R.L. Chilausky. Learning by being told and learning from examples: An experimental comparison of the two methods of knowledge acquisition in the context of developing an expert system for soybean disease diagnosis. *International Journal of Policy Analysis and Information Systems*, 4(2), 1980.

[39] M. O. Noordewier, G. G. Towell, and J. W. Shavlik. Training knowledge-based neural networks to recognize genes in DNA sequences. In R.P. Lippmann, J.E. Moody, and D.S. Touretzky, editors, *Advances in Neural Information Processing Systems-3*, pages 530–536. Morgan Kaufmann, 1991.

[40] D. W. Opitz and J. W. Shavlik. Generating accurate and diverse members of a neural-network ensemble. In D. S. Touretzky, M. C. Mozer, and M. E. Hasselmo, editors, *Advances in Neural Information Processing Systems-8*, pages 535–541. M.I.T. Press, 1996.

[41] M.P. Perrone and L. N. Cooper. Learning from what's been learned: Supervised learning in multi-neural network systems. In *Proceedings of the World Congress on Neural Networks*, pages III:354–357. INNS Press, 1993.

[42] M.P. Perrone and L. N. Cooper. When networks disagree: Ensemble methods for hybrid neural networks. In R. J. Mammone, editor, *Neural Networks for Speech and Image Processing*, chapter 10. Chapmann-Hall, 1993.

[43] Lutz Prechelt. PROBEN1 — A set of benchmarks and benchmarking rules for neural network training algorithms. Technical Report 21/94, Fakultät für Informatik, Universität Karlsruhe, D-76128 Karlsruhe, Germany, September 1994. Anonymous FTP: /pub/papers/techreports/1994/1994-21.ps.Z on ftp.ira.uka.de.

[44] J.R. Quinlan. Simplifying decision trees. *International Journal of Man-Machine Studies*, 27:221–234, December 1987.

[45] J.R. Quinlan. Bagging, boosting and C4.5. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*, Portland, OR, 1996.

[46] J. S. Rao and R. Tibshirani. Bias, variance and prediction error for classification rules. Technical report, University of Toronto, May 1997.

[47] M.D. Richard and R.P. Lippmann. Neural network classifiers estimate Bayesian a posteriori probabilities. *Neural Computation*, 3(4):461–483, 1991.

[48] G. Rogova. Combining the results of several neural network classifiers. *Neural Networks*, 7(5):777–781, 1994.

[49] D. W. Ruck, S. K. Rogers, M. E. Kabrisky, M. E. Oxley, and B. W. Suter. The multilayer Perceptron as an approximation to a Bayes optimal discriminant function. *IEEE Transactions on Neural Networks*, 1(4):296–298, 1990.

[50] A. E. Sarhan and B. G. Greenberg. Estimation of location and scale parameters by order statistics from singly and doubly censored samples. *Annals of Mathematical Statistics Science*, 27:427–451, 1956.

[51] R. Schapire. The strength of weak learnability. *Machine Learning*, 5(2):197–227, 1990.

[52] R. Schapire, Y. Freund, P. Bartlett, and Lee W.S. Boosting the margin: A new explanation for the effectiveness of voting methods. In *Proceedings of the Fourteenth International Conference on Machine Learning*. Morgan Kaufmann, 1997.

[53] A. J. J. Sharkey. (editor). *Connection Science: Special Issue on Combining Artificial Neural Networks: Ensemble Approaches*, 8(3 & 4), 1996.

[54] J. W. Smith, J. E. Everhart, W. C. Dickson, W. C. Knowler, and R. S. Johannes. Using the ADAP learning algorithm to forecast the onset of diabetes mellitus. In *Proceedings of the Symposium on Computer Applications and Medical Care*, pages 261–265. IEEE Computer Society Press, 1988.

[55] P. Sollich and A. Krogh. Learning with ensembles: How overfitting can be useful. In D. S. Touretzky, M. C. Mozer, and M. E. Hasselmo, editors, *Advances in Neural Information Processing Systems-8*, pages 190–196. M.I.T. Press, 1996.

[56] R. Tibshirani. Bias, variance and prediction error for classification rules. Technical report, University of Toronto, November 1996.

[57] K. Tumer and J. Ghosh. Order statistics combiners for neural classifiers. In *Proceedings of the World Congress on Neural Networks*, pages I:31–34, Washington D.C., 1995. INNS Press.

[58] K. Tumer and J. Ghosh. Analysis of decision boundaries in linearly combined neural classifiers. *Pattern Recognition*, 29(2):341–348, February 1996.

[59] K. Tumer and J. Ghosh. Error correlation and error reduction in ensemble classifiers. *Connection Science, Special Issue on Combining Artificial Neural Networks: Ensemble Approaches*, 8(3 & 4):385–404, 1996.

[60] D. H. Wolpert. A mathematical theory of generalization. *Complex Systems*, 4:151–200, 1990.

[61] D. H. Wolpert. Stacked generalization. *Neural Networks*, 5:241–259, 1992.

[62] D. H. Wolpert. On bias plus variance. *Neural Computation*, 9(6):1211–1243, 1997.

[63] L. Xu, M. I. Jordan, and G. E. Hinton. An alternative model for mixtures of experts. In G. Tesauro, D. S. Touretzky, and T. K. Leen, editors, *Advances in Neural Information Processing Systems-7*, pages 633–640. M.I.T. Press, 1995.

[64] L. Xu, A. Krzyzak, and C. Y. Suen. Methods of combining multiple classifiers and their applications to handwriting recognition. *IEEE Transactions on Systems, Man and Cybernetics*, 22(3):418–435, May 1992.